

# Summoning Information in Spacetime, or Where and When Can a Qubit Be?

Patrick Hayden<sup>1,2</sup> and Alex May<sup>1</sup>

<sup>1</sup>*McGill University, Montreal, Quebec, Canada*

<sup>2</sup>*Perimeter Institute for Theoretical Physics, Waterloo, Ontario, Canada*

(Dated: 12 November 2012)

While the no-cloning theorem prevents quantum information from being copied in space, the reversibility of microscopic physics *requires* that the information be copied in time. In spacetime as a whole, therefore, quantum information is widely replicated but in a restricted fashion. In this article, we fully characterize which regions of Minkowski spacetime, specifically which causal diamonds, can all hold the same quantum information. Because quantum information can be delocalized through quantum error correction and teleportation, it need not follow well-defined trajectories. Instead, replication of the information in any configuration of spacetime regions not obviously leading to violations causality and the no-cloning principle is allowed. To demonstrate this, we answer the operational question of exactly when the information can be summoned to a set of spacetime points, showing how to do so efficiently using a combination of teleportation and codeword-stabilized quantum codes.

Keywords: causality, cloning, quantum information, summoning, teleportation, error correction, relativity

To understand how information can propagate through spacetime, it is important to understand the trajectories of any possible physical information carriers. For example, a bit of classical information encoded into a beam of light at a particular point in spacetime can travel anywhere the light can. The set of possible spacetime destinations of that bit is therefore known as the point's future light cone. By copying the bit as necessary, the future light cone can essentially be filled with copies of the bit, subject only to fundamental limits on spacetime information density like the holographic bound [1].

Quantum information is very different. While a qubit's possible destinations are still restricted to lie within its future light cone, the no-cloning theorem prevents the qubit from being copied [2]. This suggests a qubit must simply follow a single physically reasonable trajectory within the light cone, which would indeed be true if the qubit were required to be localized. However, there is no such requirement. A qubit can be stored in the long-range correlations of condensed matter systems [3], delocalized using quantum error correction [4] and even teleported [5]. The purpose of this article will therefore be to understand how a quantum state can be delocalized and subsequently *relocalized*, and thereby exactly how quantum information can be distributed through spacetime.

If some quantum information is localized in the vicinity of a spacetime point, it should be possible in principle to quickly assemble and exhibit that information nearby. Indeed, that is the very definition of what it means for the information to be localized. To track propagating information, it then simply suffices to identify the extent to which the information is localized about a collection of points in spacetime. That idea can be rigorously formalized using a version of the information processing task called *summoning*, whose importance for relativistic quantum theory has already been identified by Kent [6, 7]. Suppose that some quantum information, in




FIG. 1: Spacetime diagram of a summoning task with call and reveal points separated. The state is initially localized at  $s$  and will be called at one of  $y_0$  or  $y_1$ , in which case the state must be revealed at  $z_0$  or  $z_1$ , respectively.

the form of a quantum system in an unknown quantum state  $|\varphi\rangle$ , is initially localized at some spacetime point  $s$ . A request for the state will be received at one of a set of possible spacetime *call points*  $\{y_0, y_1, \dots, y_{n-1}\}$ . Associated to each call point  $y_j$  is a *reveal point*  $z_j$  in its future light cone, with the rule that if the state is requested at  $y_j$  then it must be revealed at  $z_j$ . The more separated  $y_j$  and  $z_j$  are, the more delocalized the information can be between  $y_j$  and  $z_j$ . However, much like performing a quantum measurement, exhibiting the information at a specific  $z_j$  will eliminate all indeterminacy about its position. Therefore, it is sufficient to be able to respond to the request *in principle*. This can be enforced by stipulating that at most one of the  $n$  call points will receive the request, but which one isn't known in advance [21]. See Figure 1 for an illustration.

In characterizing the relationship between localization and summoning, it is actually possible to be more precise. Define the *causal diamond*  $D_j$  to be the intersection of the future light cone of  $y_j$  and the past light cone of  $z_j$ . This represents the region in which it is possible to both




FIG. 2: A summoning task in 1+1 dimensions. In this example, a simple strategy will work even though  $z_1$  is not in the future light cone of  $z_0$ . The quantum state is first transported along the arrow to  $P$ . The call information originating at  $y_0$  is broadcast into its future light cone and accessed at the point  $P$ . If the call is for  $z_0$ , the quantum state is moved there. If not, it is moved to  $z_1$ .

be affected by a request at  $y_j$  and affect the outcome at  $z_j$ . A qubit can be summoned to  $z_j$  from  $y_j$  if and only if the qubit is localized in the causal diamond  $D_j$ . Therefore, the summoning task is possible if and only if the qubit is localized to each and every one of the causal diamonds  $\{D_1, D_2, \dots, D_{n-1}\}$ .

To understand the definition and the surprising variety phenomena it allows, it will be helpful to consider some important examples. The simplest one consists of just two call points coinciding precisely with their associated reveal points, and in the forward light cone of a qubit [22] in the state  $|\varphi\rangle$  at the starting point. If the two call points are not spacelike to each other, then the qubit can simply be transmitted to each call point in turn. On the other hand, Kent observed that if the two call points are spacelike to each other, the impossibility of superluminal signalling implies that being able to successfully summon  $|\varphi\rangle$  would amount to being able to send  $|\varphi\rangle$  to both the reveal points, which is a clear violation of the no-cloning principle. Despite its simplicity, this no-summoning theorem has significant consequences for information processing. In the same way that the no-cloning theorem gives rise to secret key distribution protocols secured by the laws of quantum mechanics, the no-summoning theorem gives rise to secure bit commitment protocols secured by a combination of quantum mechanics and relativity [8]. Secure bit commitment is impossible using quantum mechanics alone [9, 10].

Another simple example in 1+1 dimensions is shown in Figure 2. Even though there is no causal curve through each of the reveal points, it is still possible to complete the summoning task. The example illustrates that if there is a causal curve passing through each of the causal diamonds, then summoning will be possible.




FIG. 3: Exploiting quantum error correction. One share of a  $((2, 3))$  threshold quantum secret sharing scheme is allocated to each of the call points  $y_j$ . Meanwhile, each reveal point  $z_j$  is lightlike to both  $y_j$  and  $y_{j+1 \bmod 3}$ . (The vertical direction is time. The arrows are causal, in this case lightlike, curves.) If the participants follow the protocol described in the text, then the correct reveal point will receive enough shares to reconstruct the secret.

When more spatial dimensions are introduced, delocalization becomes crucial in understanding which configurations of causal diamonds are compatible. For a representative example, place three call points at time zero on the vertices of an equilateral triangle with edge lengths  $\ell$ . Place the reveal points at time  $\ell/(2c)$  on the midpoints of the edges, for  $c$  the speed of light. Because the call-reveal pairs are lightlike, the diamonds  $D_i$  are just line segments, as shown in Figure 3. There is no causal curve through the diamonds  $D_i$ , so the strategy of simply moving the qubit around won't work. Delocalizing the qubit through the use of quantum error correcting codes will, however. It is possible to encode the quantum state  $|\varphi\rangle$  into a tripartite Hilbert space  $\mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \mathcal{H}_3$  such that the qubit can be recovered even if any one of the corresponding three quantum subsystems is lost. This is known as a  $((2, 3))$  threshold quantum secret sharing scheme because the quantum information can be recovered from any two of the three subsystems even though no information at all can be recovered from fewer than two [11]. This encoding is performed at the start point  $s$  and then one share is forwarded to each of the call points. For each  $j$ , if the request is made at  $y_j$ , then that share is forwarded to  $z_j$ . Otherwise, that share is forwarded to  $z_{j-1 \bmod 3}$ . By this arrangement the correct reveal point will receive two out of the three shares as required to recover the state.

The example reveals one interesting way to delocalize a quantum state. To address the full variety of ways in which quantum information can be distributed in space-time, the following theorem characterizes every summoning task in Minkowski space as possible or impossible.

**Theorem 1** *Summoning is possible if and only if the following conditions hold:*

1. Every reveal point is in the future light cone of the

starting point  $s$ .

2. For each pair  $(i, j)$ , the diamonds  $D_i$  and  $D_j$  are causally related, meaning that there exists a causal curve from  $D_i$  to  $D_j$  or vice versa.

The two conditions are necessary because they encode the most basic constraints coming from relativity and quantum mechanics, namely causality and the impossibility of cloning. Indeed, Condition 1 is manifestly the prohibition of superluminal communication. Condition 2 arises from reasoning similar to Kent's treatment of "non-ideal" summoning [6]. Suppose we have a successful summoning protocol for which Condition 2 is violated, meaning that two diamonds  $D_i$  and  $D_j$  are spacelike separated as in Figure 1. If the call is received at  $y_i$ , there is a procedure that will reveal the state at  $z_i$ . Now imagine that the call machinery malfunctions such that it makes a call at  $y_j$  in addition to the one at  $y_i$ . Because  $y_j$  is not in the causal past of  $z_i$ , the malfunction cannot prevent the state from being revealed at  $z_i$ . Likewise, because  $y_i$  is not in the causal past of  $z_j$ , the call at  $y_j$  will result in the state successfully being revealed at  $z_j$ . This procedure therefore reveals the state  $|\varphi\rangle$  at the two spacelike points  $z_i$  and  $z_j$  starting from a single copy of  $|\varphi\rangle$  at the point  $s$ . In other words, a summoning protocol for a configuration violating Condition 2 is easily modified to make a cloning machine, which is impossible.

To see that Conditions 1 and 2 are sufficient will require constructing a protocol that will succeed at the summoning task given a starting point and  $n$  call-reveal pairs satisfying the conditions. The structure of the protocol will only depend on the directed graph  $G = (V, E)$  whose vertices are labelled by the diamonds  $D_i$  and which contains the edge  $(D_i, D_j)$  if and only if there is a causal curve from some point in  $D_i$  to one in  $D_j$ .

It is possible to handle the  $n = 2$  case by making use of a strategy from [7]. Without loss of generality, assume there is a causal curve from  $D_0$  to  $D_1$ . Begin by distributing a Bell pair between the spatial locations of the start point and  $y_0$ . Upon receiving the quantum state at the start point, immediately teleport it over the Bell pair [5], sending the classical teleportation data to both  $z_0$  and  $z_1$ . Meanwhile, if the call is received at  $y_0$ , forward the other half of the Bell pair to  $z_0$ , but if no call is received, forward it to  $z_1$ . Because there is a causal curve from the start point to both  $z_0$  and  $z_1$ , and because there is a causal curve from  $D_0$  and  $D_1$  (which, in particular, guarantees there is a causal curve from  $y_0$  to  $z_1$ ), both the classical data and the half of the Bell pair required to reconstruct the quantum state will arrive at the appropriate reveal point. Figure 4 depicts an example in which this protocol succeeds but the simpler strategy of carrying the qubit through the causal diamonds fails.

Using quantum error correction, a protocol for general  $n$  can be built recursively from the protocol for  $n = 2$ . Encode the state  $|\varphi\rangle$  at the starting point  $s$  in an  $((n - 1, n))$  threshold secret sharing scheme [11]. There are  $n$  subsets of  $\{D_1, D_2, \dots, D_n\}$  of size  $n - 1$ .




FIG. 4: The general strategy for completing two-request tasks can be used to complete this example, even though  $y_0$  and  $y_1$  are outside the light cone of  $s$ . The essence of teleportation is that it splits a qubit into entanglement and classical data transmission, thereby making it possible to delocalize quantum information in a curious way: classical data can be transmitted to several recipients without regard to the no-cloning theorem while entanglement reaches outside the light cone. Both features are crucial in this example.




FIG. 5: a) This configuration of four call-reveal spacetime points satisfies the conditions of Theorem 1 and requires a combination of error correction and teleportation for successful summoning. Each reveal point is again lightlike to its call point and  $z_j$  is causal to the call point  $y_{j-1 \bmod 4}$ . In addition,  $z_0$  is lightlike to  $y_2$  and  $z_3$  is lightlike to  $y_1$ . b) The graph  $G$  of causal relationships used to construct the summoning protocol. The subproblem involving the cycle  $D_0 \rightarrow D_2 \rightarrow D_3 \rightarrow D_0$  is structurally equivalent to the example illustrated in Figure 3, although the recursive protocol replaces direct quantum communication by teleportation. c) The code used in the efficient construction has one physical qubit for each edge of this graph,  $G'$ . The arcs label the four sets of doubled edges sufficient to reconstruct the state. In the protocol, if there is a request at  $y_0$ , for example, the qubits on the solid edges are sent to  $z_0$ .

Assign one of the  $n$  shares to each of the subsets and for each subset recursively execute the protocol, now on the smaller subset of size  $n - 1$ . If the request is made at call point  $y_j$ , then for each of the subsets containing  $D_j$ , the corresponding protocol will forward its share of the secret to  $z_j$ . Precisely  $n - 1$  of the  $n$  subsets contain  $D_j$ , so the state  $|\varphi\rangle$  will be recoverable at  $z_j$ , as required. An example for  $n = 4$  is sketched in Figures 5a and 5b.

**Efficient construction**— The protocol described in

the proof of Theorem 1 is unfortunately inefficient, using  $\Omega(n!)$  qubits. Practicality aside, such dramatic growth quickly runs afoul of the holographic bound, which places a limit of roughly  $1.4 \times 10^{69}$  bits per  $\text{m}^2$  [1]: trying to store the protocol's qubits in a region centred at  $s$  of area of  $1 \text{ m}^2$  would already create a black hole for  $n = 55$ . Understanding summoning in the presence of gravity therefore requires finding more efficient protocols.

The high cost is incurred from the recursive encoding preceding teleportation, so we will show how to achieve the same functionality directly. To begin, it will be helpful to characterize that functionality. The recursive encoding produces a quantum error correcting code with shares associated to the edges of the graph  $G$ . Should the call arrive at  $y_j$ , teleportation is used to ensure that all shares associated to the edges incident to  $D_j$  arrive at the reveal point  $z_j$ . The functionality required of the code is therefore clear: it must be possible to recover the encoded quantum state using only those shares received at  $z_j$ . That is, the erasure of shares for any combination of edges leaving at least one vertex untouched must be recoverable, as illustrated in Figure 5c.

That describes an unusual quantum error correcting code. Ignoring directionality,  $G$  is the complete graph, so each vertex is incident to exactly  $n-1$  edges. The total number of edges is  $\binom{n}{2}$ , however, so the quantum state must be recoverable from a vanishing fraction  $2/n$  of the total number of shares, albeit a specially chosen vanishing fraction that prevents violations of the no-cloning principle. A particular codeword-stabilized (CWS) quantum code [12] will do the job very efficiently. Let  $G' = (V', E')$  be an undirected version of  $G$  with a new vertex placed on each edge, so that  $V' = V \cup E$  and  $E' = \{(v, (v, w)) : (v, w) \in E \text{ or } (w, v) \in E\}$ . The code will consist of a single qubit for each of the edges of  $G'$ . For each such edge  $e$ , let  $N_e = \{f \in E' : f \cap e \neq \emptyset \text{ and } f \neq e\}$  be the set of edges adjacent to  $e$  and define  $S_e = X_e \prod_{f \in N_e} Z_f$ , where  $X_e$  and  $Z_e$  are Pauli operators acting on edge  $e$ . The encoded qubit is simply the span of the simultaneous  $+1$  eigenstate of all the  $S_e$  operators and the simultaneous  $-1$  eigenstate.

As demonstrated in Appendix B, this code has precisely the desired properties if the share for each edge  $e \in E$  of the original graph  $G$  is identified with the pair of edges replacing  $e$  in  $G'$ . The construction therefore yields a method for solving the summoning task using exactly  $n(n-1)$  physical qubits per summoned qubit. The CWS code also happens to be a stabilizer code, so the erasure errors can be decoded in polynomial time [4].

**Generalizations**— The summoning tasks characterized by Theorem 1 only require that it be possible to respond to a request at a single call point. That version of the problem is the one most relevant to understanding how information can be distributed through spacetime. From an active information processing perspective, however, it is natural to permit requests at *sequences* of call points rather than just one. These sequential summoning tasks will be specified by a set  $\{\bar{y}_0, \bar{y}_1, \dots, \bar{y}_{n-1}\}$  of




FIG. 6: A summoning task defined by sequences of call-reveal pairs; the three possible sequences are indicated by dashed arrows through the reveal points.

possible sequences of call points and a corresponding set  $\{\bar{z}_0, \bar{z}_1, \dots, \bar{z}_{n-1}\}$  of sequences of reveal points. Any sequence of requests  $\bar{y}_j$  must be met by revealing the state at each of the paired reveal points. A typical problem is illustrated in Figure 6.

The characterization of which sequential summoning tasks are possible reduces easily to Theorem 1. It is clearly necessary that there be a causal curve through each set of reveal points  $\bar{z}_j$ . If not, then there would be a sequence containing two spacelike separated reveal points, and to complete the task would violate no-signalling. Without loss of generality then, the sequences  $\bar{y}_j$  and  $\bar{z}_j$  can be assumed to be in the order induced by the causal curve through  $\bar{z}_j$ . Use the first element of each sequence to define a summoning subtask. Since only one sequence of call points will receive requests, only one of the call points of this smaller task will, and Theorem 1 will apply. At completion of this first task, the quantum state will be localized at the requested reveal point. Take this as a new start point and consider the summoning subtask consisting of the second element of each sequence whose first reveal point is this new start point. All other sequences can be safely discarded because they are not consistent with the first request. Because there is a causal curve from the first element in each sequence to the second, this second subtask will have each of its reveal points in the future light cone of the start point. Thus Theorem 1 applies and the task can be completed. Repeating this process and completing each of the subtasks in turn completes the multiple request task. This analysis implies that sequential summoning is possible if and only if there is causal curve through each  $\bar{z}_j$  and the conditions of Theorem 1 hold for each subtask.

**Discussion**— Our analysis of summoning and its sequence-based generalization gives a detailed characterization of the ways in which quantum information can be distributed through spacetime. It is remarkable that the achievable combinations of information requests can be so cleanly characterized; the conditions of Theorem 1 are simple constraints arising from causality and no-cloning,

but they prove to be sufficient. Thus, the most basic restrictions imposed by quantum mechanics and special relativity turn out to be the only restrictions.

These conclusions should extend to more general spacetimes since our arguments have only depended on causal structure. Ultimately, however, the reasoning will likely to break down in a full theory of quantum gravity, for which it has been argued that the interplay between cloning and causality should be much more subtle [13, 14]. Most importantly, the use of no-cloning to rule out summoning for a pair of spacelike separated causal diamonds likely fails in some situations in which the structure of spacetime prevents the two clones from ever being compared.

A question closely related to the problem studied here is whether quantum mechanics and relativity can be combined to build a “position-based” cryptography, which was recently answered in the negative [15, 16]. In both that work and ours, novel combinations of entanglement and teleportation make it possible to transmit quantum information in ways that initially seem to defy causality. Nonetheless, both scenarios prove to be quite amenable to analysis, and can be seen as building blocks [7, 17–20] for a larger theory of distributed computation in relativistic spacetime.

*Acknowledgments*— We are indebted to Adrian Kent, Fang Xi Lin and Prakash Panangaden for helpful discussions. This work was supported by the Canada Research Chairs program, the Perimeter Institute, CIFAR, FQRNT’s INTRIQ, NSERC and ONR through grant N000140811249. The Perimeter Institute is supported by Industry Canada and Ontario’s Ministry of Economic Development & Innovation.

## Appendix A: On definitions and assumptions

There are a number of assumptions implicit in the definition of summoning supplied in the main text. Kent does an excellent job of clearly identifying them in [6], to which we refer the reader for a more thorough discussion. Most importantly, we assume that information at any spacetime point can be transmitted without errors to any other point in its forward light cone, and that quantum computations at any point can be performed instantaneously.

Those assumptions are necessary in order to be able to assert that a qubit is localized to the causal diamond  $D_j$  if and only if it can be summoned to  $z_j$ . In a realistic physical theory, features of the dynamics beyond causal structure and quantum mechanics could prevent the summons from succeeding. The quantum state of planet Earth exactly one day ago, for example, is localized within a spatial sphere one light-day in radius, but it doubtful that the state could be reassembled here on Earth one day hence.

Summoning with a separation between the call and reveal point as we’ve defined it in this article is actually a

version of “non-ideal summoning” in Kent’s terminology, which he analyzed for a case in which all the call points are on one spacelike hyperplane and all the reveal points on another. Kent introduced the separation in order to allow time to process information between the receipt of the request and the need to reveal the state. Because our emphasis has been on the ways quantum information can be localized in principle, our interpretation of the separations has been slightly different, but nothing prevents using them for Kent’s original purpose.

## Appendix B: Analysis of the error correcting code

The CWS formalism constructs a quantum error correcting code from an undirected graph  $H$  and a classical error correcting code, associating a physical qubit with each vertex of  $H$  [12]. Since we require a share for each *edge* of  $G$ , we will choose  $H$  to be the line graph of the undirected graph  $G'$ . The classical error correcting code will simply be a length  $2\binom{n}{2}$  binary repetition code. These choices uniquely specify the code described in the main text. Because the binary repetition code forms a group, this CWS code is also a stabilizer code [12].

The conditions for quantum error correction discovered in [12] can easily be re-expressed in terms of the graph  $G'$ . CWS codes map all Pauli errors into patterns of  $Z$  errors. In particular, the error  $X_e$  error on any edge  $e$  is mapped into the product  $\prod_{f \in N_e} Z_f$  of  $Z$  errors on all the edges adjacent to  $e$ . Up to an irrelevant sign, therefore, any Pauli error  $P$  induces a product of  $Z$  errors  $\text{err}(P)$ .

As discussed in the main text, it must be possible to correct against the erasure of any shares leaving any vertex  $u$  of  $G = (V, E)$  untouched. That is equivalent to being able to correct any set of Pauli errors located on edges of  $G'$  not associated with  $u$ . Since the edges  $E'$  of  $G'$  are of the form  $(v, (v, w))$  for  $v \in V$  and  $(v, w) \in E$ , that amounts to be able to correct errors on any  $(v, (v, w))$  such that  $u \notin (v, w)$ . Let  $\mathcal{E}$  be the set of such errors for any fixed  $u$ . Since our construction is invariant with respect to permutations of the vertices of  $V$ , the choice of  $u$  doesn’t matter.

The error correction conditions from [12] specialized to our case are that, for each  $P \in \mathcal{E}$ ,

- $\text{err}(P) \neq \prod_e Z_e$  and
- either  $\text{err}(P) \neq I$  or  $[\prod_e Z_e, P] = 0$ .

The first condition is easy to verify. Let  $P \in \text{err}(P)$  be a single-qubit error acting on edge  $(v, (v, w))$ . Then  $u \notin (v, (v, w))$ . The only edge adjacent to  $(v, (v, w))$  containing  $u$  is  $(v, (v, u))$ . In particular, there are no edges of the form  $(u, (u, x))$  adjacent to  $(v, (v, w))$  so  $P$  cannot induce an error on any such edge, and neither can a product of single-qubit errors in  $\mathcal{E}$ . Therefore  $\text{err}(P) \neq \prod_e Z_e$  for all  $P \in \mathcal{E}$ .

For the second condition, begin by considering a single-qubit  $X$  error  $P \in \mathcal{E}$ , again acting on edge  $(v, (v, w))$ .

$(v, (v, u))$  is the only edge adjacent to  $(v, (v, w))$  containing  $u$ , so  $\text{err}(P)$  will contain exactly one  $Z$  operator acting on edges containing  $u$ . Using that  $Z^2 = I$ , it follows that any multiqubit  $X$  error  $P \in \mathcal{E}$  for which  $\text{err}(P) = I$  must contain an even number of single-qubit errors. But

$XZ = -ZX$  so  $[\prod_e Z_e, P] = 0$  for any such error  $P$ . A general multiqubit error will contain arbitrary products of  $X$  and  $Z$  errors, but since  $Z$  errors commute with  $\prod_e Z_e$  and don't propagate along edges, the second condition holds for them as well.

- 
- [1] R. Bousso. The holographic principle. *Rev. Mod. Phys.*, 74:825–874, 2002.
  - [2] W.K. Wootters and W.H. Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, 1982.
  - [3] A. Y. Kitaev. Fault-tolerant quantum computation by anyons. *Ann. Phys.*, 303:2–30, 2003.
  - [4] D. Gottesman. *Stabilizer codes and quantum error correction*. PhD thesis, California Institute of Technology, 1997.
  - [5] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters. Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Phys. Rev. Lett.*, 70:1895–1899, 1993.
  - [6] A. Kent. A no-summoning theorem in relativistic quantum theory. *Quant. Inf. Proc.*, pages 1–10, 2011.
  - [7] A. Kent. Quantum tasks in Minkowski space. *Class. Quantum Grav.*, 29:224013, 2012.
  - [8] A. Kent. Unconditionally secure bit commitment with flying qudits. *New J. Phys.*, 13(11):113015, 2011.
  - [9] D. Mayers. Unconditionally secure quantum bit commitment is impossible. *Phys. Rev. Lett.*, 78:3414–3417, 1997.
  - [10] H.-K. Lo and H. F. Chau. Is quantum bit commitment really possible? *Phys. Rev. Lett.*, 78:3410–3413, 1997.
  - [11] R. Cleve, D. Gottesman, and H.-K. Lo. How to share a quantum secret. *Phys. Rev. Lett.*, 83:648–651, 1999.
  - [12] A. Cross, G. Smith, J.A. Smolin, and B. Zeng. Codeword stabilized quantum codes. In *IEEE International Symposium on Information Theory, 2008. ISIT 2008.*, pages 364–368. IEEE, 2008.
  - [13] L. Susskind, L. Thorlacius, and J. Uglum. The stretched horizon and black hole complementarity. *Phys. Rev. D*, 48(8):3743, 1993.
  - [14] P. Hayden and J. Preskill. Black holes as mirrors: quantum information in random subsystems. *JHEP*, 09:120, 2007.
  - [15] A. Kent, W. J. Munro, and T. P. Spiller. Quantum tagging: Authenticating location via quantum information and relativistic signalling constraints. *Phys. Rev. A*, 84(1):012326, 2011.
  - [16] Harry Buhrman, Nishanth Chandran, Serge Fehr, Ran Gelles, Vipul Goyal, Rafail Ostrovsky, and Christian Schaffner. Position-based quantum cryptography: Impossibility and constructions. In *Advances in Cryptology: CRYPTO 2011*, volume 6841 of *Lecture Notes in Computer Science*, pages 429–446. Springer, 2011.
  - [17] D. Beckman, D. Gottesman, M. A. Nielsen, and J. Preskill. Causal and localizable quantum operations. *Phys. Rev. A*, 64(5):052309, 2001.
  - [18] D. Beckman, D. Gottesman, A. Kitaev, and J. Preskill. Measurability of Wilson loop operators. *Phys. Rev. D*, 65(6):065022, 2002.
  - [19] M. Pawłowski, T. Paterek, D. Kaszlikowski, V. Scarani, A. Winter, and M. Żukowski. Information causality as a physical principle. *Nature*, 461(7267):1101–1104, 2009.
  - [20] S. Beigi and R. König. Simplified instantaneous non-local quantum computation with applications to position-based cryptography. *New J. Phys.*, 13(9):093036, 2011.
  - [21] A more extended discussion of the assumptions implicit in this definition can be found in Appendix A.
  - [22] For linguistic convenience, we always discuss summoning a qubit, but our results extend trivially to any finite-dimensional quantum information.